B.J.Pol.S. **50**, 381–389
doi:10.1017/S0007123417000412
First published online 1 November 2017

Note

A Get-Out-the-Vote Experiment on the World's Largest Participatory Budgeting Vote in Brazil

TIAGO PEIXOTO, FREDRIK M. SJOBERG AND JONATHAN MELLON*

Keywords: get out the vote; experiment; turnout; vote choice; participatory budgeting; Brazil

While academic studies of voter mobilization have taught us much about what drives turnout,[1] we know little about the effect that get-out-the-vote (GOTV) campaigns have on election outcomes. Experimental work has shown that there are heterogeneous treatment effects – in other words, that encouragements have different turnout effects on different people.[2] However, there is little evidence from turnout experiments about whether mobilization efforts affect the outcome of the vote. One reason the GOTV literature primarily focuses on turnout is data availability. Individual vote choice is difficult to observe directly due to the secret ballot, while administrative turnout records can be obtained in many places. The few previous studies of GOTV on vote choice look at partisan mobilization and rely on incomplete self-reported or aggregate vote choice.[3] As a result, we do not know whether GOTV campaigns actually mobilize people who vote differently than those who would otherwise have voted.

Several studies have looked at the relationship between turnout and policy outcomes at the national level, and consistently find that higher turnout is associated with higher levels of redistribution.[4] However, these studies rely on observational data and cannot assign higher turnout experimentally.

This study has two main outcomes of interest. First, we examine the vote encouragement treatment effects on online turnout in a participatory budgeting (PB) vote. This part of the study is similar to much of the established GOTV literature, in which the outcome of interest is whether an eligible voter turns out to vote on election day. Secondly, we study the effect of GOTV treatments on the actual vote using administrative records of vote choice. To our knowledge, this has never been studied before using real vote data. The nature of the PB vote in our case allows us to study both the average cost of proposals selected by voters as well as specific sectoral (health, education, environment etc.) preferences.

* Digital Engagement Evaluation Team, World Bank (emails: tpeixoto@worldbank.org, fredrik.m.sjoberg@gmail.com), Digital Engagement Evaluation Team, World Bank and Manchester University (email: jonathan.mellon@manchester.ac.uk). The pre-analysis plan was submitted to the Evidence in Governance and Politics pre-registration repository prior to random assignment and treatment application (20140523AA). We would like to thank Vincius Wu from the Rio Grande do Sul Government, Motta, Davi Schmidt, Paulo Coelho at SEPLAG, Uirá Porã, Luiz Damasceno and the rest of the staff at Gabinete Digital, Rosane Maria Ludtke Leite and Guilherme Donato at PROCERGS, and Louis Dorval at Voto Mobile. Funding for this research was provided by the World Bank. Note that no pre-approval by an Institutional Review Board was sought for this study, since no such process exists within the World Bank. However, the study was approved *post facto* by three World Bank research staff who were not involved with the original research; they agreed that it adhered to the research and ethical standards of World Bank research. Data replication sets are available in Harvard Dataverse at: https://dx.doi.org/10.7910/DVN/PSWLYR, and an online appendix at: https://doi.org/10.1017/S0007123417000412.

[1] Gerber and Green 2000; Gerber, Green, and Larimer 2008; Green, McGrath, and Aronow 2013.
[2] Enos, Fowler, and Vavreck 2014; Imai and Strauss 2010.
[3] Arceneaux 2005; Arceneaux 2007; Pons 2014; Rogers and Middleton 2015.
[4] Filer, Kenny, and Morton 1993; Larcinese 2007; Mahler 2008.

To test whether increased turnout from voter mobilization affects the popular vote, we conducted a randomized controlled trial during the June 2014 participatory budgeting vote in Rio Grande do Sul (RS) in southern Brazil.[5] Note that this is not a regular election in which votes are cast for candidates and parties. In a PB vote, voters are asked to choose between specific public investment projects that they can indicate their support for. We devised a procedure in close collaboration with the implementing government agency that retained an indicator of individual-level treatment assignment with the actual vote choice record. The procedure anonymizes all other individual attributes, thus preserving the secrecy of the ballot.

The list of experimental subjects consists of 43,384 voters who took part in the annual online PB vote in the past two years. The list includes only individual voters who voluntarily provided both an email address and a phone number when voting in the past. The population we study appears to have a higher propensity to vote than the rest of the online voting population. Examining historical online PB vote data, we find that 22.4 per cent of voters for whom we have an email address took part in the online vote in both 2012 and 2013, compared to 18.2 per cent of all online voters.

We randomly assigned subjects, in equal proportions, to one of four groups: control, informational, public benefit of voting, and private benefit (lottery reward) using a simple random allocation scheme.[6] All of the messages were non-partisan in nature and were focused on increasing turnout. The sender of the email was *Gabinete Digital*, the governor's digital engagement unit. Figure 1 shows the email that subjects in the informational treatment group received in the morning on the first day of online voting (2 June). Each experimental subject was sent three emails: a voter registration email on 30 May, a message about voting starting on 2 June, and a message about it being the last chance to vote on 4 June. We also sent out a mobile phone text message (SMS) on 3 June.

The content of the message was the same throughout the process for each treatment condition, while the exact wording of SMS messages varied due to space limitations. On 2 June, the first treatment group was sent an informational message stating basic information about this year's online voting taking place from 2–4 June. The message subject line read 'Today: Priorities Vote!' and the body text contained a direct link to the voting site. The two other groups were also sent the informational message, but the public benefit motivation condition was also assigned the following message: 'This is your chance to make government work on your priorities.' In addition, the subject line for the public benefit condition included the words: 'Make a difference.' The private benefit condition instead received the following message: 'If you vote in via Citizen Login you can win two extra tickets for the raffle at R$ 1 million occurring in March 2015.'

In this GOTV experiment, we study the final stage of an annual three-month long budget consultation process that culminates in citizens voting on their priorities for a pre-screened set of public investment projects and policy sectors. In the PB vote in RS the ballot is split into two sections: a list of up to 20 possible public investment projects with specific costs attached to each of them, of which a voter can pick up to four projects, and a list of five possible regional sectoral (health, education, etc.) priorities, of which a voter can choose up to two options. The items on the ballot vary across the twenty-eight PB areas depending on the projects suggested within the PB meetings held prior to voting day, and a few of the regions do not list costs on the ballot.[7] However, on all ballots, each of these items is coded into one of the same fourteen possible thematic areas, allowing us to compare vote choices across these electoral areas.

Because the secrecy of the vote is a key tenet of electoral integrity, the GOTV literature has had to rely on self-reported or aggregate-level voting when studying vote choice. To ensure that individual voters could not be identified in our study, the state agency that administered the online vote (PROCERGS) merged a vector with the treatment assignment into the vote records data prior to decrypting an individual's vote choice. After the vote verification process was completed (checking for duplicate voters, voters voting in the wrong area, etc.), the vote record stripped individual identifiers, while retaining the treatment assignment vector. Only at this point were the votes decrypted. The vote records could then be transferred to the research team safely, without compromising individuals' privacy. At no point in the voting process

[5] For more information about the PB process in Rio Grande do Sul, see Spada et al. (2016).

[6] The lottery reward was a pre-existing voter encouragement scheme that all voters were eligible for the lottery message merely provided information about the scheme.

[7] The Appendix contains a sample ballot.

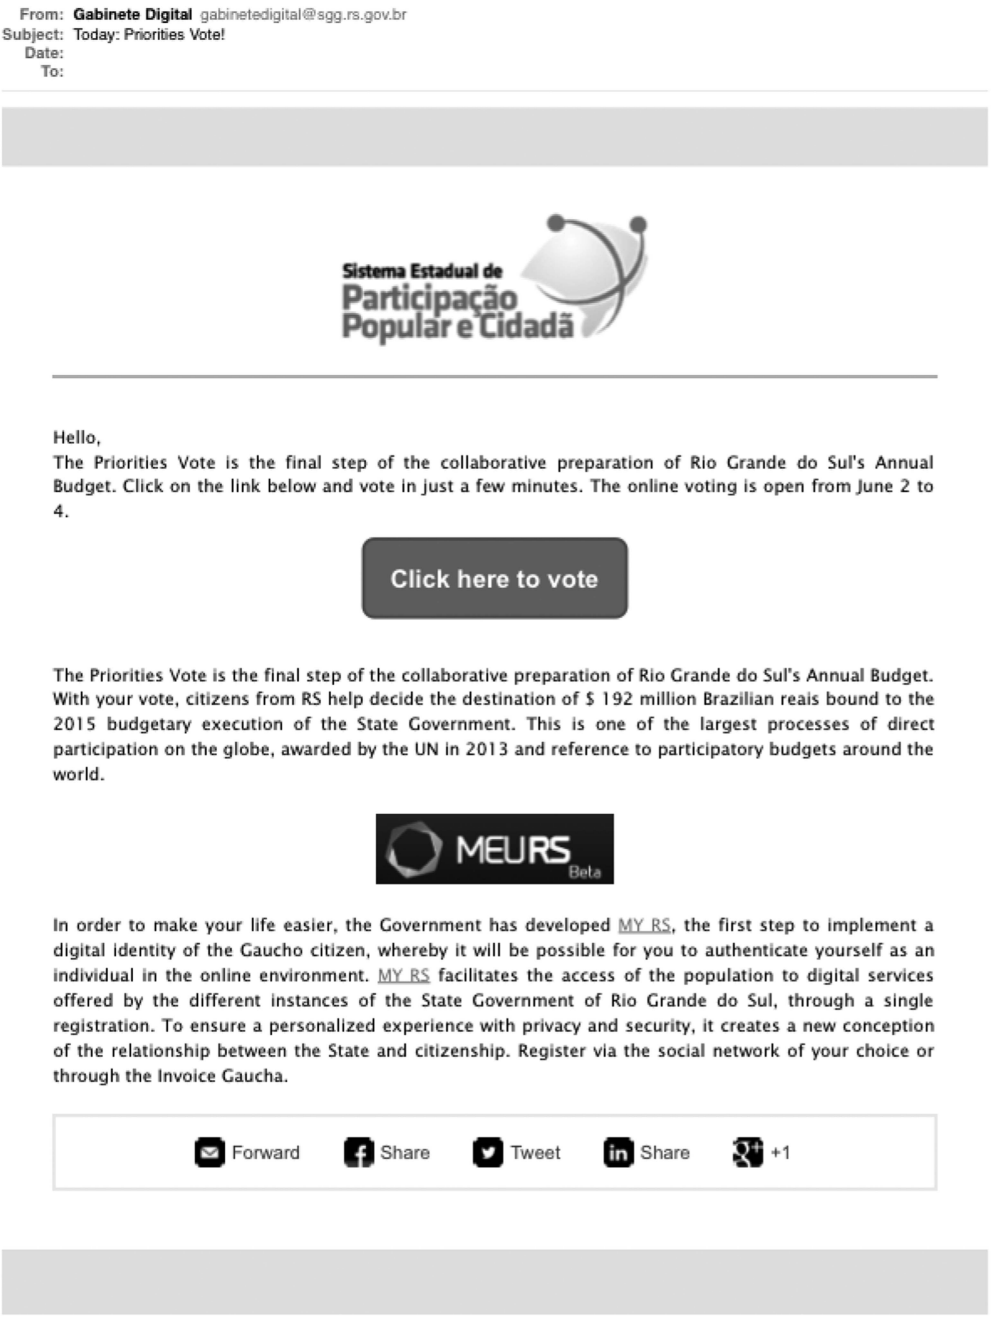

Fig. 1. Informational treatment message on election day 1 on June 2
Note: more details about the treatments are available in the Appendix.

were decrypted votes and personally identifiable information kept on the same server. This approach is analogous to voters being given a different colored ballot when their identification is checked at the polling station depending on their treatment group. The analysis is then conducted solely in terms of differences in votes cast using different colored ballots. Since the number of experimental groups is

https://doi.org/10.1017/S0007123417000412 Published online by Cambridge University Press

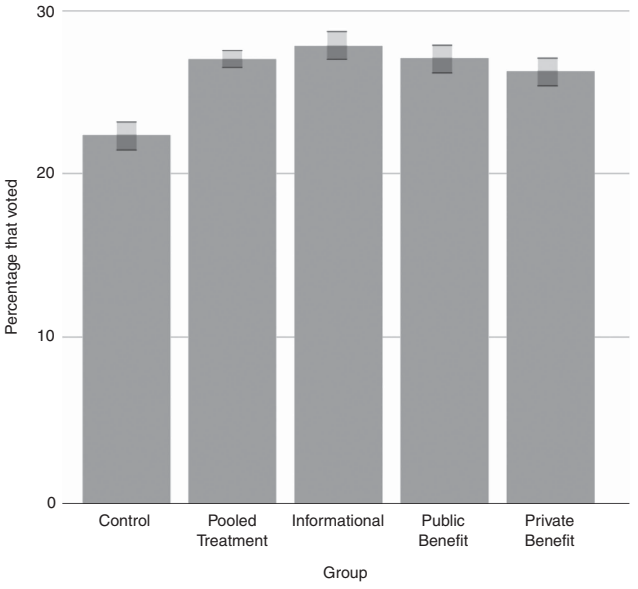

Fig. 2. Online turnout in control, pooled treatment group and the three treatments separately
Note: the three treatments are informational, public benefit and private benefit. 95 per cent confidence intervals are displayed.

extremely small compared to the number of voters, the different ballots do not meaningfully increase the *post hoc* identifiability of voters.

Prior to treatment assignment, we submitted a pre-analysis plan to Experiments in Governance and Politics. As specified in the plan, we estimate the intent-to-treat (ITT) effect by using a difference-in-means estimator. Compliance with treatment, defined as the recipient having read the subject line or the body text, is difficult to determine since we do not know whether the mail ended up in a recipient's junk folder, but the bounce rate was a very low 1.6 per cent. With non-compliance this low, effectively the estimand is the Average Treatment Effect (ATE). In the linear model we include dummy variables for each of the twenty-eight regions (COREDES) since the ballot varies across regions. We address the multiple comparisons problem using Benjamini and Hochberg's method for adjusting p-values[8] to control the false discovery rate (FDR) for each dependent variable.[9]

The first substantive finding is that different GOTV treatments have a turnout-increasing effect (see Figure 2). Taken together, voters assigned to receive any of the GOTV messages were 4.7 percentage points more likely to vote in the online poll (ITT, difference in means linear estimation, $p < 0.001$, see Table 1). Turnout in the control group was 22.6 per cent. Each of the separate mobilization treatments shows significant differences from the control group that received no email. The content of the mobilization messages does not matter much, a finding that is broadly consistent with the literature.[10] However, the private benefit (lottery) treatment did perform significantly worse than the informational treatment (−1.6 percentage points, difference in means linear estimation, FDR-adjusted $p = 0.01$).

The turnout effect we document here is substantially larger than that found in most previous GOTV studies,[11] particularly those focusing on the effect of new technologies, like email and SMS.[12]

[8] Benjamini and Hochberg 1995.
[9] Sun et al. 2006.
[10] Gerber and Green 2000.
[11] Green, McGrath, and Aronow 2013; Nickerson 2007.
[12] Bhatti et al. 2015; Malhotra et al. 2011; Nickerson 2007.

https://doi.org/10.1017/S0007123417000412 Published online by Cambridge University Press

TABLE 1 *Treatment Effects on Turnout – Main Model and Comparisons Across Different Treatment Subsets*

Comparison	N	Estimate	Std. Error	Unadjusted p-value	FDR adjusted p-value
Pooled treatments	43,384	0.047	0.005	0.000	0.000
Information	43,384	0.055	0.006	0.000	0.000
Public benefit	43,384	0.047	0.006	0.000	0.000
Private benefit	43,384	0.039	0.006	0.000	0.000
Info vs. Public	21,692	−0.009	0.006	0.158	0.185
Public vs. Private	21,692	−0.008	0.006	0.209	0.209
Info vs. Private	21,692	−0.016	0.006	0.008	0.011

While we cannot test the reasons for this performance directly, it is likely due to repeated messages and multi-mode contact, and the ability to vote by clicking on a link in the message. In addition, while this is a low-salience election, almost a quarter of the subjects targeted were habitual voters, for whom large effects have been found.[13] The high performance of the informational message suggests that one of the primary ways in which the GOTV message is working is simply by informing voters that there is an election. The lower performance of the private benefit message could be for a number of reasons. For example, a private benefit may undermine public benefit motivation,[14] it may be more likely to be caught by spam filters, or voters may simply be less trusting of such a message. Regardless of the reason, these results suggest that large lottery rewards are not an especially effective way of mobilizing voters.

In the second part of the study, we examined unique administrative vote records. Previous studies have shown that non-voters differ from voters,[15] which implies that additional voters brought in by GOTV efforts will have an effect on election outcomes by altering the preferences represented. We test this implication directly by examining vote choice across the different treatment groups. As specified in the pre-analysis plan, we first test whether our treatments affect the average cost of proposals selected by voters across different groups. Each proposal was listed on the ballot with a Brazilian Real (R$) cost attached to it. While the total amount of money spent is not directly affected by the choices a voter makes, we view a lower total cost on a ballot as reflecting a voter's preference for lower spending. We operationalize cost as the average cost of the items a voter chose on the ballot, expressed as a proportion of the maximum cost they could have picked on their ballot. In the control group that number was 65.1 per cent. As can be seen from Table 2, the total costs do not differ significantly between the control and treatment groups (–0.2 percentage points, difference in means linear estimation, FDR-adjusted $p = 0.79$). Similarly, none of the individual treatments shows differences from the control group or each other. These results show that encouraged voters do not differ from natural voters in terms of their preferences for higher- or lower-costing public investments.

In addition to the cost of proposals, we also consider whether encouraged and un-encouraged voters differ in their preferences for different public policy sectors. We compare the proportion of voters that voted for an item in the three most popular issue areas displayed on the second part of the ballot. In the 2014 PB vote, these were health, infrastructure and security. We also include an 'other' category for the proportion that voted for any of the other categories. We analyze only the votes on the second section of the ballot, where a voter can choose up to two priorities for their own region (*Prioridades Regionais*).

Table 3 shows the proportions that chose each area across the control and treatment conditions. As with the total cost, we find no differences between the control and pooled treatment groups in their preferences for health (FDR-adjusted $p = 0.98$), infrastructure (FDR-adjusted $p = 0.98$), security (FDR-adjusted $p = 0.98$) or other spending (FDR-adjusted $p = 0.98$). We also tested whether voters in the different treatment groups differed from each other if, for instance, voters encouraged to participate by a public

[13] Malhotra et al. 2011.
[14] Fehr and Falk 2002; Fehr and Fischbacher 2002.
[15] Enos, Fowler, and Vavreck 2014.

https://doi.org/10.1017/S0007123417000412 Published online by Cambridge University Press

TABLE 2 *Treatment Effects on Cost*

Comparison	N	Estimate	Std. Error	Unadjusted p-value	FDR adjusted p-value
Pooled treatments	43,384	−0.002	0.006	0.794	0.794
Information	43,384	−0.008	0.007	0.263	0.717
Public benefit	43,384	−0.003	0.007	0.662	0.773
Private benefit	43,384	0.007	0.007	0.320	0.717
Info vs. Public	21,692	0.004	0.007	0.512	0.717
Public vs. Private	21,692	0.010	0.007	0.132	0.717
Info vs. Private	21,692	0.004	0.007	0.497	0.717

Note: main model and comparisons across different treatment subsets. Average cost of options chosen on a respondent's ballot as a percentage of the maximum cost they could have chosen.

TABLE 3 *Treatment Effects on Vote Choice*

Treatment	Issue Area	N	Estimate	Std. Error	Unadjusted p-value	FDR adjusted p-value
Pooled	Health	11,350	−0.003	0.011	0.775	0.984
Pooled	Infrastructure	11,350	0.006	0.011	0.608	0.984
Pooled	Security	11,350	−0.003	0.009	0.731	0.984
Pooled	Other	11,350	−0.008	0.008	0.278	0.984
Information	Health	11,350	−0.012	0.014	0.379	0.984
Public	Health	11,350	0.002	0.014	0.911	0.984
Private	Health	11,350	0.001	0.014	0.943	0.984
Information	Infrastructure	11,350	0.015	0.013	0.263	0.984
Public	Infrastructure	11,350	0.002	0.013	0.902	0.984
Private	Infrastructure	11,350	0.000	0.014	0.984	0.984
Information	Security	11,350	−0.005	0.011	0.643	0.984
Public	Security	11,350	−0.007	0.011	0.529	0.984
Private	Security	11,350	0.003	0.011	0.806	0.984
Information	Other	11,350	−0.007	0.009	0.457	0.984
Public	Other	11,350	−0.010	0.009	0.267	0.984
Private	Other	11,350	−0.008	0.009	0.384	0.984
Info vs. Public	Health	6,015	0.014	0.013	0.276	0.984
Info vs. Public	Infrastructure	6,015	−0.013	0.013	0.322	0.984
Info vs. Public	Security	6,015	−0.002	0.010	0.828	0.984
Info vs. Public	Other	6,015	−0.003	0.009	0.721	0.984
Public vs. Private	Health	5,840	−0.001	0.013	0.949	0.984
Public vs. Private	Infrastructure	5,840	−0.002	0.013	0.902	0.984
Public vs. Private	Security	5,840	0.010	0.011	0.356	0.984
Public vs. Private	Other	5,840	0.002	0.009	0.807	0.984
Info vs. Private	Health	8,894	0.014	0.013	0.268	0.984
Info vs. Private	Infrastructure	8,894	−0.014	0.013	0.282	0.984
Info vs. Private	Security	8,894	−0.002	0.010	0.842	0.984
Info vs. Private	Other	8,894	−0.003	0.009	0.723	0.984

Note: main model and comparisons across different treatment subsets. The average proportion of voters choosing the top three sectoral categories: health, infrastructure or security, and an 'other' category combining all other categories.

benefit message might be more inclined to support health projects. In all cases, we found no differences in sectoral preferences according to treatment condition. Taken together, these results provide no evidence that our GOTV treatments had any impact on the outcome of the online vote.

Based on the size of the GOTV effect and the proportions of voters in the treatment and control groups, we can examine the implied size of the difference in preferences between regular (voters who would have voted anyway) and encouraged (voters who only voted because of the GOTV treatment) voters

TABLE 4 *Estimated Differences in Sectoral Preference Between Regular and Encouraged Voters and Difference in Chosen Cost*

Issue	Estimate	Lower CI	Upper CI
Infrastructure	0.041	−0.084	0.164
Health	−0.018	−0.127	0.080
Security	−0.026	−0.134	0.036
Other	−0.048	−0.145	0.116
Proposal cost	−0.015	−0.083	0.057

Note: confidence intervals estimated using bootstrapping.

(although none of these differences is close to significant, as we discussed above). Applying the method outlined by Fowler,[16] the largest implied difference is voting for 'other', which we estimate as 4.8 percentage points less likely among encouraged voters (see Table 4). However, all of these differences have confidence intervals encompassing zero. In the Appendix, we show a power analysis that suggests that we would expect to be able to detect a 19-percentage-point difference 83 per cent of the time. This means that we can be reasonably confident that there are no large differences in sectoral preferences between regular and encouraged voters. We cannot, for instance, rule out the possibility that differences in the size seen here of 1.8–4.8 percentage points might remain and become significant with a larger sample size.

An important implication of our study is that email messaging can be a powerful tool to mobilize people to take part in a PB vote, especially if the type of engagement allows for online modes of participation. This is particularly vital in the context of Brazil, which has experienced a wave of protests related to, among other grievances, government spending. Democratic innovations that allow citizens to participate directly in budgeting are seen as an important way of increasing citizen engagement and government legitimacy, and high rates of participation are considered very important for the credibility and sustainability of these initiatives.[17]

Future research should focus on replicating these results in other contexts. As the world moves toward more online voting, analyzing vote choice may become more feasible in more cases. However, we would encourage researchers to think carefully about voter privacy. The focus in this study, as in most of the literature, was on non-partisan voter mobilization. Needless to say, we still expect that partisan GOTV campaigns targeting their own core supporters or particular demographic groups would shift the election outcome by increasing their own supporters' turnout. However, this would be the result of differential targeting rather than heterogeneous treatment effects.

This research calls into question the assumption that changing the pool of voters by increasing turnout will inevitably affect the distribution of votes cast. Unlike observational studies, we do not find that higher turnout affects voters' policy preferences, and we can be confident that the new voters brought in via the GOTV campaign do not have drastically different sectoral preferences. There are a number of possible reasons for these divergent findings. One possibility is that another factor explains both higher turnout and redistribution across countries. However, there are also important contextual differences that could explain our findings. In a low-salience and non-partisan election, the voters who are mobilized through GOTV may be sufficiently similar to those who were already turning out to vote that voters' choices are not distinguishable. It is also possible that endogenous increases in turnout are more likely to bring in voters who vote differently to existing voters than increases in turnout generated through untargeted GOTV. Finally, the choices in a PB vote may be less correlated with turnout than in partisan elections, which have more established signals for voters to use. Nonetheless, it is striking that the experimental evidence of turnout's effect on policy choices comes to such a different conclusion than previous observational evidence.

[16] Fowler 2015.
[17] Goldfrank and Schneider 2006.

https://doi.org/10.1017/S0007123417000412 Published online by Cambridge University Press

While it might initially seem that our results imply that there is no value in using non-partisan GOTV to increase turnout, this assumes that the only value of voting is to change the outcome of an election. Many arguments have been made for increasing turnout, including legitimizing the political system, decreasing the alienation of the populace[18] and even as a form of education.[19] The results of this study should encourage non-partisan GOTV campaigns precisely because they are unlikely to greatly affect the results.

REFERENCES

Arceneaux, Kevin. 2005. Using Cluster Randomized Field Experiments to Study Voting Behavior. *The Annals of the American Academy of Political and Social Science* 601:169–79.

——. 2007. I'm Asking for Your Support: The Effects of Personally Delivered Campaign Messages on Voting Decisions and Opinion Formation. *Quarterly Journal of Political Science* 2:43–65.

Benjamini, Yoav, and Yosef Hochberg. 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57:289–300.

Bhatti, Yosef, Jens Olav Dahlgaard, Jonas H. Hansen, and Kasper M. Hansen. 2015. Getting Out the Vote With Evaluative Thinking. *American Journal of Evaluation* 36:389–400.

Enos, Ryan D., Anthony Fowler, and Lynn Vavreck. 2014. Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate. *The Journal of Politics* 76:273–88.

Fehr, Ernst, and Armin Falk. 2002. Psychological Foundations of Incentives. *European Economic Review* 46:687–724.

Fehr, Ernst, and Urs Fischbacher. 2002. Why Social Preferences Matter–The Impact of Non-Selfish Motives on Competition, Cooperation and Incentives. *The Economic Journal* 112:C1–33.

Filer, John E., Lawrence W. Kenny, and Rebecca B. Morton. 1993. Redistribution, Income, and Voting. *American Journal of Political Science* 37:63–87.

Fowler, Anthony. 2015. Regular Voters, Marginal Voters and the Electoral Effects of Turnout. *Political Science Research and Methods* 3:205–19.

Gerber, Alan S., and Donald P. Green. 2000. The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment. *American Political Science Review* 94:653–63.

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. Social Pressure and Voter Turnout: Evidence From a Large-Scale Field Experiment. *American Political Science Review* 102:33–48.

Goldfrank, Benjamin, and Aaron Schneider. 2006. Competitive Institution Building: The PT and Participatory Budgeting in Rio Grande Do Sul. *Latin American Politics & Society* 48:1–31.

Green, Donald P., Mary C. McGrath, and Peter M. Aronow. 2013. Field Experiments and the Study of Voter Turnout. *Journal of Elections, Public Opinion & Parties* 23:27–48.

Imai, Kosuke, and Aaron Strauss. 2010. Estimation of Heterogeneous Treatment Effects From Randomized Experiments, With Application to the Optimal Planning of the Get-Out-the-Vote Campaign. *Political Analysis* 19:1–19.

Larcinese, Valentino. 2007. Voting Over Redistribution and the Size of the Welfare State: The Role of Turnout. *Political Studies* 55:568–85.

Mahler, Vincent A. 2008. Electoral Turnout and Income Redistribution by the State: A Cross-National Analysis of the Developed Democracies. *European Journal of Political Research* 47:161–83.

Malhotra, Neil, Melissa R. Michelson, Todd Rogers, and Ali Adam Valenzuela. 2011. Text Messages as Mobilization Tools: The Conditional Effect of Habitual Voting and Election Salience. *American Politics Research* 39:664–81.

Nickerson, David W. 2007. The Ineffectiveness of E-Vites to Democracy Field Experiments Testing the Role of E-Mail on Voter Turnout. *Social Science Computer Review* 25:494–503.

Pateman, Carole. 1970. *Participation and Democratic Theory*. Cambridge: Cambridge University Press.

[18] Powell 1986; Warren 1999.
[19] Pateman 1970.

Peixoto, Tiago, Fredrik M. Sjoberg, and Jonathan Mellon. 2017. Replication Data for: A Get Out The Vote (GOTV) Experiment on the World's Largest Participatory Budgeting Vote in Brazil. doi:10.7910/DVN/SWLYR, Harvard Dataverse, V1, UNF:6:w3WwhWET1CyJgTqa2osvYg==.

Pons, Vincent. 2014. Does Door-To-Door Canvassing Affect Vote Shares? Evidence from a Countrywide Field Experiment in France.

Powell Jr, G. Bingham. 1986. American Voter Turnout in Comparative Perspective. *The American Political Science Review* 80:17–43.

Rogers, Todd, and Joel Middleton. 2015. Are Ballot Initiative Outcomes Influenced by the Campaigns of Independent Groups? A Precinct-Randomized Field Experiment Showing That They Are. *Political Behavior* 37:567–93.

Spada, Paolo, Jonathan Mellon, Tiago Peixoto, and Fredrik Matias Sjoberg. 2016. Effects of the Internet on Participation: Study of a Public Policy Referendum in Brazil. *Journal of Information Technology and Politics* 13:187–207.

Sun, Lei, Radu V. Craiu, Andrew D. Paterson, and Shelley B. Bull. 2006. Stratified False Discovery Control for Large-Scale Hypothesis Testing With Application to Genome-Wide Association Studies. *Genetic Epidemiology* 30:519–30.

Warren, Mark E. 1999. *Democracy and Trust*. Cambridge: Cambridge University Press.